Mon. Not. R. Astron. Soc. **000**, 000–000 (0000) Printed 4 March 2016 (MN LATEX style file v2.2)

Machine-z: Rapid Machine Learned Redshift Indicator for *Swift* Gamma-ray Bursts

T. N. Ukwatta[1⋆], P. R. Woźniak[2] and N. Gehrels[3]
[1] *Director's Postdoctoral Fellow, Space and Remote Sensing (ISR-2), Los Alamos National Laboratory, Los Alamos, NM 87544, USA.*
[2] *Space and Remote Sensing (ISR-2), Los Alamos National Laboratory, Los Alamos, NM 87544, USA.*
[3] *Astroparticle Physics Division, NASA/Goddard Space Flight Center, Greenbelt, MD 20771, USA.*

arXiv:1512.07671v2 [astro-ph.HE] 3 Mar 2016

4 March 2016

ABSTRACT

Studies of high-redshift gamma-ray bursts (GRBs) provide important information about the early Universe such as the rates of stellar collapsars and mergers, the metallicity content, constraints on the re-ionization period, and probes of the Hubble expansion. Rapid selection of high-z candidates from GRB samples reported in real time by dedicated space missions such as *Swift* is the key to identifying the most distant bursts before the optical afterglow becomes too dim to warrant a good spectrum. Here we introduce "machine-z", a redshift prediction algorithm and a "high-z" classifier for *Swift* GRBs based on machine learning. Our method relies exclusively on canonical data commonly available within the first few hours after the GRB trigger. Using a sample of 284 bursts with measured redshifts, we trained a randomized ensemble of decision trees (random forest) to perform both regression and classification. Cross-validated performance studies show that the correlation coefficient between machine-z predictions and the true redshift is nearly 0.6. At the same time our high-z classifier can achieve 80% recall of true high-redshift bursts, while incurring a false positive rate of 20%. With 40% false positive rate the classifier can achieve ∼100% recall. The most reliable selection of high-redshift GRBs is obtained by combining predictions from both the high-z classifier and the machine-z regressor.

Key words: gamma-ray bursts, redshift

1 INTRODUCTION

Gamma-ray bursts (GRBs) are often characterized as the most energetic electromagnetic explosions since the beginning of the Universe. Their optical afterglows are in principle detectable out to redshift $z > 10$ (Lamb and Reichart 2000; Mesler et al. 2014). Therefore, studies of distant GRBs can probe the physics of the early Universe including the re-ionization, the evolution of star formation, and the process of metal enrichment (Lamb and Reichart 2000; Totani et al. 2006; Kawai et al. 2006). High-redshift GRBs can be used to pinpoint and characterize the faint galaxies that supplied most of the re-ionization photons and to constrain the re-ionization redshift (Ioka 2003; Totani et al. 2006; Wang et al. 2012). Multi-wavelength studies of distant GRB afterglows can provide new information about the metal and dust content of these objects (Frail et al. 2006; Cusumano et al. 2006; Mesler et al. 2014). This in turn offers a unique method to understand the metal enrichment history of sources during the re-ionization epoch. While most studies of the re-ionization epoch are based on observations of quasars, high-z GRBs have a number of significant advantages over quasars due to their unique characteristics. GRB afterglows are exceptionally bright and provide plenty of photons for sensitive spectroscopy. They have simple, easy to model power-law spectra dominated by the continuum emission that are well suited for detecting absorption signatures of the intergalactic medium. Additionally, the neighborhoods of GRB progenitors are relatively "clean" compared to quasars, which are often contaminated by continuous ejection of material from the central engine.

The *Swift* Gamma-Ray Burst Mission (Gehrels et al. 2004) has proven to be effective in detecting very high-redshift GRBs. The most distant spectroscopically confirmed GRB on record is GRB 090423 with $z = 8.2$ (Tanvir et al. 2009; Salvaterra et al. 2009)[1]. The highest photometrically measured burst is GRB 090429B with a redshift of 9.4 (Cucchiara et al. 2011). There is now a handful of spectroscopically confirmed GRBs with $z > 5$. The main chal-

⋆ E-mail: tilan@lanl.gov; tilan.ukwatta@gmail.com

[1] GRB 120923A may have slightly higher redshift. A preliminary analysis of the photometric data indicates $z \sim 8.5$ (Tanvir 2013).

lenge in this work is to reliably identify high-redshift bursts suitable for detailed spectroscopic follow-up before the optical emission fades away. Decisions to use precious observing time on large telescopes must be made within the first hours or even minutes after the burst based on limited information. Previous attempts to screen high-z GRBs using promptly available high-energy data (Campana et al. 2007; Salvaterra et al. 2007; Koen 2009, 2010; Ukwatta et al. 2008, 2009; Morgan et al. 2012), while showing some promise, lacked the accuracy necessary to facilitate a reliable follow-up program. As a result, they were never widely adopted by observers.

The main difficulty lies in extracting numerous weak correlations from readily available high-dimensional data and efficiently combining the information they contain. A modern approach based on machine learning is ideal for this purpose. Starting from a catalog of GRBs with known redshifts we can use supervised learning to "train" algorithms that effectively encode the relationship between input data and output labels. Classification algorithms deal with predicting discrete labels (in this case high-z versus low-z), while regression algorithms predict continuous labels (here the redshift value). Both types of models are supported by the random forest algorithm that in recent years has emerged as one of the best performing machine learning tools in observational astrophysics.

In this paper we present a rapid machine learned redshift estimator called *machine-z* and a *high-z* classifier for GRBs detected by *Swift*. Both *machine-z* and *high-z* are developed independently and each tool may be used to reinforce conclusions from the other. In Section 2 we describe the input data and the method. Our high-z classifier is developed in Section 3 and the machine-z indicator is developed in Section 4. In Section 5 we compare our algorithms and results with previous work and evaluate the performance of our new tools using a sample of recently detected bursts that are not included in the training catalog. In Section 6 we summarize the results.

2 METHODOLOGY

2.1 GRB Sample

Our sample consists of 284 *Swift* GRBs with spectroscopic redshift measurements[2]. The *Swift* mission payload consists of three major instruments: the Burst Alert Telescope (BAT), the X-ray Telescope (XRT) and the UV Optical Telescope (UVOT) (Gehrels et al. 2004). BAT is a soft gamma-ray wide field instrument sensitive to photons in the energy range 15 keV to 350 keV and it is the GRB discovery instrument. Once BAT discovered a GRB and determined its sky position, the *Swift* satellite slews to the location of the burst so that the narrow field instruments XRT and UVOT can quickly start observing the afterglow. In order to provide a rapid *machine-z* redshift and *high-z* classification, we limited this study to readily available measurements from all three *Swift* instruments. These measurements were adopted as numerical features for classification and regression, and are listed in Table 1.

[2] The measurements are taken from the *Swift* online catalog at `http://swift.gsfc.nasa.gov/archive/grb_table/`.

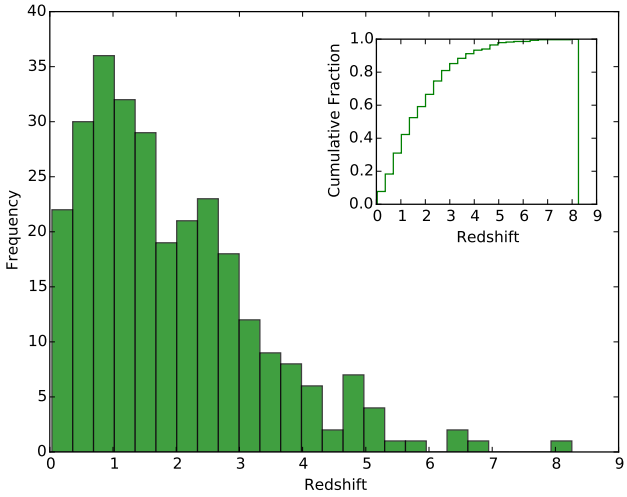

Figure 1. Redshift distribution for a sample of 284 *Swift* bursts.

In total we considered 25 features. The features generally capture the timing and spectral properties of the prompt and afterglow emission from the burst. Some measurements are encoded as two separate features to ensure that all available information is included. For example, the prompt emission recorded by BAT is fitted using either a power law (PL) model or a cutoff power law (CPL) model, depending on which one provides a better fit. The power-law index is an important feature, but the existence of a high-energy cutoff provides additional information and is included as a separate feature `FitType` that encodes the best fit model (0 for PL and 1 for CPL). Similarly, UVOT magnitudes are reported as either detections or upper limits. A numerical flag similar to `FitType` is introduced to include this information for each photometric band of UVOT.

The sample includes GRBs discovered between 2005 and the end of 2014. Fig. 1 shows the redshfit distribution of all bursts in the sample. The lowest redshift value is 0.033 and the highest is 8.26. We adopt $z = 4$ as the threshold between low-redshift and high-redshift bursts. Out of 284 GRBs in the sample 25 or $\sim 9\%$ are high-redshift according to this definition.

2.2 Machine Learning Algorithm

2.2.1 Random Forests

The random forest (RF) algorithm (Breiman 2001) has been shown to provide superior performance on many classification problems (Caruana and Niculescu-Mizil 2006; du Buisson et al. 2015; D'Isanto et al. 2016). Over the past few years the method found several interesting applications in observational astrophysics including selection of explosive transients in imaging data (Wright et al. 2015; Goldstein et al. 2015), classification of X-ray sources (Farrell et al. 2015), and redshift prediction (Carliles et al. 2010; Morgan et al. 2012). RF has the ability to select useful features, relatively immune to data over-fitting, can handle nonlinear relationships, and provide probabilistic outputs (see Morgan et al. (2012) and references therein). Given input training data the algorithm creates a large number of decorrelated binary

Table 1. Standard features used for training classification and regression algorithms. The features are derived from measurements by all three *Swift* instruments: BAT, XRT and UVOT.

Item	Feature	Units	Instrument
1	T90	sec	BAT
2	Fluence (15-150 keV)	10^{-7}erg/cm^2	BAT
3	1-sec Peak Photon Flux (15-150 keV)	ph/cm^2/sec	BAT
4	Photon Index	None	BAT
5	Fit Type (CPL - Cutoff Power Law, PL - Power Law)	CPL=1, PL=0	BAT
6	Early Flux (0.3-10 keV)	10^{-11}erg/cm^2/s	XRT
7	11 Hour Flux (0.3-10 keV)	10^{-11}erg/cm^2/s	XRT
8	24 Hour Flux (0.3-10 keV)	10^{-11}erg/cm^2/s	XRT
9	Initial Temporal Index	None	XRT
10	Spectral Index (Γ)	None	XRT
11	Column Density (NH)	10^{21}cm^{-2}	XRT
12	V Mag/Limit Value	Magnitudes	UVOT
13	V Value Type	Mag=1, Limit=0	UVOT
14	B Mag/Limit Value	Magnitudes	UVOT
15	B Value Type	Mag=1, Limit=0	UVOT
16	U Mag/Limit Value	Magnitudes	UVOT
17	U Value Type	Mag=1, Limit=0	UVOT
18	UVW1 Mag/Limit Value	Magnitudes	UVOT
19	UVW1 Value Type	Mag=1, Limit=0	UVOT
20	UVM2 Mag/Limit Value	Magnitudes	UVOT
21	UVM2 Value Type	Mag=1, Limit=0	UVOT
22	UVW2 Mag/Limit Value	Magnitudes	UVOT
23	UVW2 Value Type	Mag=1, Limit=0	UVOT
24	White Mag/Limit Value	Magnitudes	UVOT
25	White Value Type	Mag=1, Limit=0	UVOT

decision trees. Each tree in the forest is grown by splitting the portion of the training data associated with a particular parent node between two child nodes according to the value of one or more features. Splits are chosen to maximize node purity (classification) or minimize variance (regression). The process starts from the root node that holds the entire data set and continues until the size of the leaf node falls below a specified threshold. Randomness enters in two distinct ways. A new bootstrap sub-sample is drawn from the original training sample to train the next tree. Then a random subset of all available features is selected to optimize each split. The size of the random subset is specified by the user.

Prediction for a single tree is accomplished by propagating a previously unseen feature vector starting from the root and until a leaf node is reached. The predicted label is typically a majority class of the leaf node (classification) or a numerical average (regression). The results from all trees in the forest are then combined to compute the posterior probability of possible outcomes and the final prediction. Typically the majority vote is adopted for classification and the mean for regression. All results described in this paper were obtained using the python implementation of RF distributed with the scikit-learn package[3] (Pedregosa et al. 2012).

2.2.2 Missing Features

Missing features are common in real world data and our GRB sample is no exception. A popular approach to handle missing input values is by imputation i.e. assigning values estimated from the distribution of all remaining instances. The method works well unless missing values carry a special meaning in a given particular problem domain. In our input GRB catalog some features could not be extracted due to low signal to noise ratio of the original data or a non-detection in a particular photometric band. Both occurrences are actually expected to be correlated with redshift (distance) and are therefore examples of informative missing features. Blue filter "dropouts" are especially interesting because they are often indicative of high redshift. In order to preserve all available information about the redshift we assigned all missing features to -1000. The effect on performance is negligible as long as the plug-in value is well outside the normal range for all features. Decision trees are generally very good in utilizing such special values if the missing features are in fact informative or marginalizing them out if they are not.

2.2.3 Data Imbalance

Another common problem when searching for rare objects of interest is highly uneven distribution of training data between classes. The redshift distribution in Fig. 1 represents competition between survey volume growing rapidly with distance and decreasing efficiency of detecting more distant bursts. The input catalog for our study is quite unbalanced with fewer than 10% of GRBs at $z > 4$. A low fraction of high-redshift bursts in the training sample will typically result in a tendency to classify all bursts as low-redshift as the algorithm attempts to minimize the overall error rate. This imbalance will result in poor performance of the classifier on new data. One possible solution is to assign higher weights to high-z bursts during training. However, the price is often additional complexity and a tendency for overfitting. It is

[3] `http://scikit-learn.org`

Table 2. Parameter grid used to approximately optimize algorithm learning.

Parameter	Values
ntrees	50, 100, 200, 300, 400, 500
nodesize	1, 2, 5, 8, 10, 12, 15, 18, 20
m	2, 5, 8, 10, 12, 15, 18, 20, 22, 25

not clear at this point what is the optimal way to introduce weights in random forest and almost certainly the answer depends on the problem at hand. A preliminary investigation of the redshift bias discussed in Section 4.3 indicates that the underlying cause is imbalanced training data combined with noisy features in a significant fraction of the sample. A detailed treatment of these intricate effects warrants a separate investigation and will be presented elsewhere.

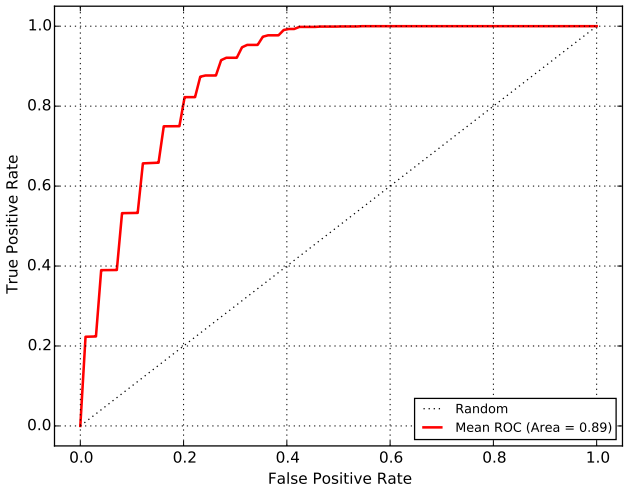

Figure 3. ROC curve for *high-z* classifier using the best 8 features identified in Fig. 2.

3 HIGH-Z CLASSIFICATION

3.1 Receiver Operating Characteristic (ROC) Curve

The receiver operating characteristic (ROC) curve is a convenient way to track performance and compare classifiers and/or feature sets. We compute this curve using randomized 10-fold cross-validation (train on 90% of the sample and test on 10%). A single point on the ROC curve corresponds to an average of 100 independent cross-validation runs for a fixed threshold applied to the probability of the burst having a high redshift. The curve is traced by varying the threshold between 0 and 1 (see Fig. 3 for an example). A perfect classifier has zero false positive rate and 100% efficiency (upper left corner of the diagram). Fast rising ROC curves are generally better than a slow rising ones. The area under the curve can be used as a rough measure of classification performance. An ideal classifier has the area of 1, while completely random selection on average yields half of the total area of the diagram.

3.2 Tuning the Classifier

RF classifiers take several parameters that can be tuned to improve performance. Among those the most important are the number of trees in the forest (*ntrees*), the minimum size of the leaf node (*nodesize*), and the number of randomly selected features that will be used to optimize node splits (m). In order to approximately optimize the *high-z* classifier we performed a simple parameter search over a grid given in Table 2 and selected a set of parameters with the largest area under the ROC curve. This results in a forest of 300 trees with at least 12 training samples per node and $m = 25$ random features per split that delivers an ROC curve with an area of 0.87. The fact that $m = 25$ is preferred means that the best results are obtained with a large degree of randomness injected during construction of individual trees (all available features are randomized). The area under the ROC curve is somewhat insensitive to the exact combination of parameters.

3.3 Classification Feature Importance

While RF is relatively immune to correlated and uninformative features, it is still beneficial to investigate the relative importance of input features on performance. We start with a pool of available features that initially contains all features in Table 1. The first most informative feature is selected to maximize classification performance (area under ROC curve) using only one feature at a time from the pool of N available features. The next best feature is selected after looping over $N - 1$ features remaining in the pool and maximizing classification performance using two features. The process continues until the pool of available features is empty. We use parameter values from Section 3.2, i.e. $ntrees = 300$ and $nodesize = 12$, except m which cannot be larger than the number of features selected for a given iteration. The relative importance of all 25 classification features is shown in Figure 2. As more features are included in training, the area under the ROC curve increases rapidly with a maximum value of 0.89 around 8-th feature followed by a gradual decrease. It is interesting to note that the 8 best features selected in this way include information from all three *Swift* instruments. The absence of the total burst duration (BAT T90) on this list is somewhat surprising and may be attributed to a large intrinsic spread of burst durations that dominates the influence of time dilation on this time scale.

The ROC curve of the final *high-z* classifier trained using the 8 best features is shown on Fig. 3. The curve shows a steep rise and reaches 100% recall at about 40% false posive rate. We can reduce the false positive rate by half by changing the probability threshold and accepting 80% recall.

3.4 Machine Learned Scoring

The ROC curve is not the only way to measure the performance of selecting high-redshift GRBs. Morgan et al. (2012) introduced a scoring method suitable for scheduling follow-up observations constrained by limited resources (telescope time). The idea is to compare *high-z* probabilities of all bursts in the training sample to that of the new event under

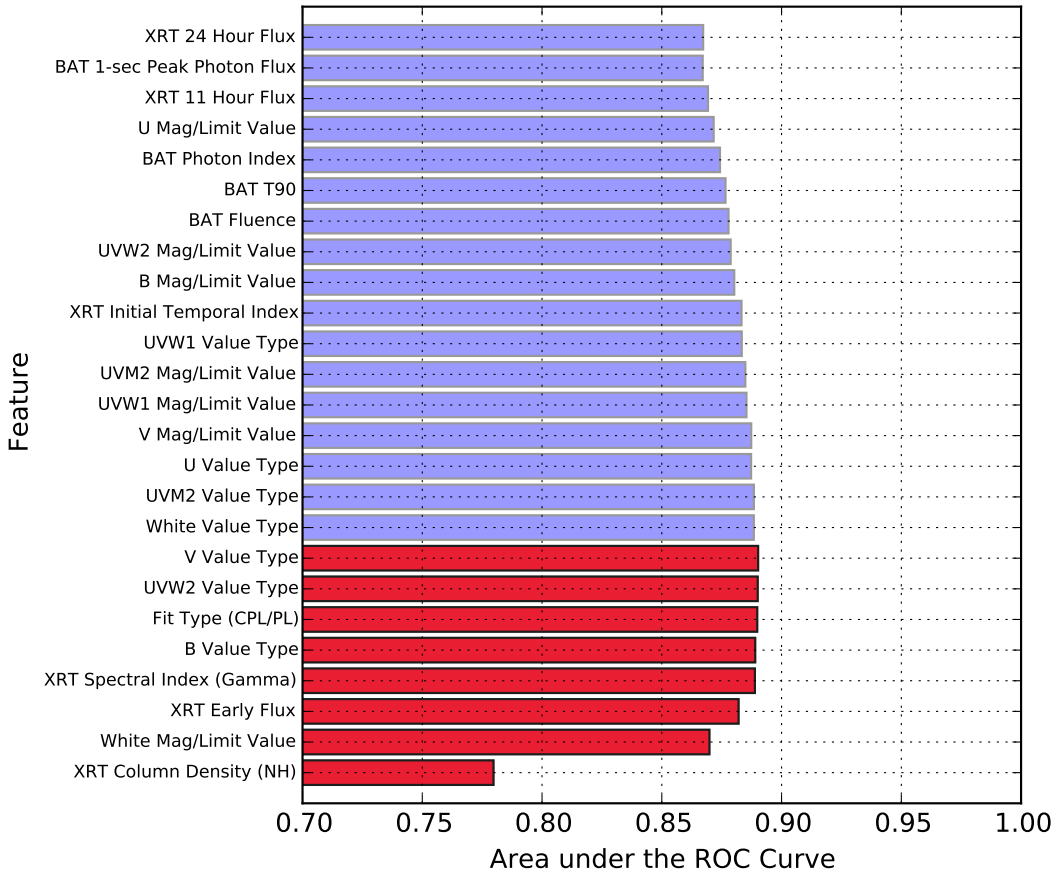

Figure 2. Relative importance of *high-z* classification features. The area under the ROC curve is shown as a function of the next best feature starting from the single best feature at the bottom of the plot. Features selected for the final *high-z* classifier are shown in red.

consideration. The new event is assigned a rank n, with $n-1$ previously seen events having a higher probability of being high-z compared to the new event. The learned follow-up rank of the new event is $Q = n/(N+1)$. This formulation supports "go or no go" decisions for newly detected bursts, given the fraction F of all GRBs that can be followed up with the currently available resources. If $Q < F$ the event should be observed. Otherwise it makes more sense to wait for a better candidate. The algorithm automatically adapts to changes in the redshift distribution of the input GRB sample and the amount of telescope time that can be devoted to follow-up.

We tested the above approach by simulating follow-up decisions for bursts on our training sample using a randomized cross-validation procedure described in Section 3. The Q scores were calculated using approximately optimized input parameters and features found in Sections 3.2 and 3.3. The effectiveness of this hypothetical observing campaign as a function of the requested follow-up fraction F is shown in Fig. 4. The top panel (a) shows that the fraction of bursts recommended for follow-up by the algorithm ($Q < F$) closely tracks the requested value F. The middle panel (b) shows the purity of the follow-up sample, i.e. the number of actual high-z GRBs divided by the total number of selected high-z candidates. Ideally, the purity would be close to 100% when the follow-up resources are limited (low F), as shown by the green line. The bottom panel (c) shows the efficiency of selecting high-z GRBs (the fraction of all high-z bursts that were actually observed). Again, perfect classification performance is shown by the green line.

From Fig. 4 it is clear that the classifier can identify high-redshift bursts with a high probability, especially when follow-up resources are very limited (low F). In other words, the purity is highest when very few bursts can be followed up and therefore reliable predictions matter most. At 1% follow-up fraction the purity exceeds 80%. On the other hand, an observer with enough telescope time to follow up 40% of all GRBs will be able to find all true high-z bursts (100% efficiency).

4 MACHINE-Z REDSHIFT ESTIMATOR

The high-z classifier developed in section 3 helps to select the highest priority follow-up targets, but it does not provide an actual value for the predicted redshift. In this section, we adapt the methods from section 3 to solve a regression problem and develop an RF based redshift estimator that we call *machine-z*.

4.1 Tuning the Regressor

Since the performance of a regressor is evaluated differently from a classifier, we performed an independent parameter

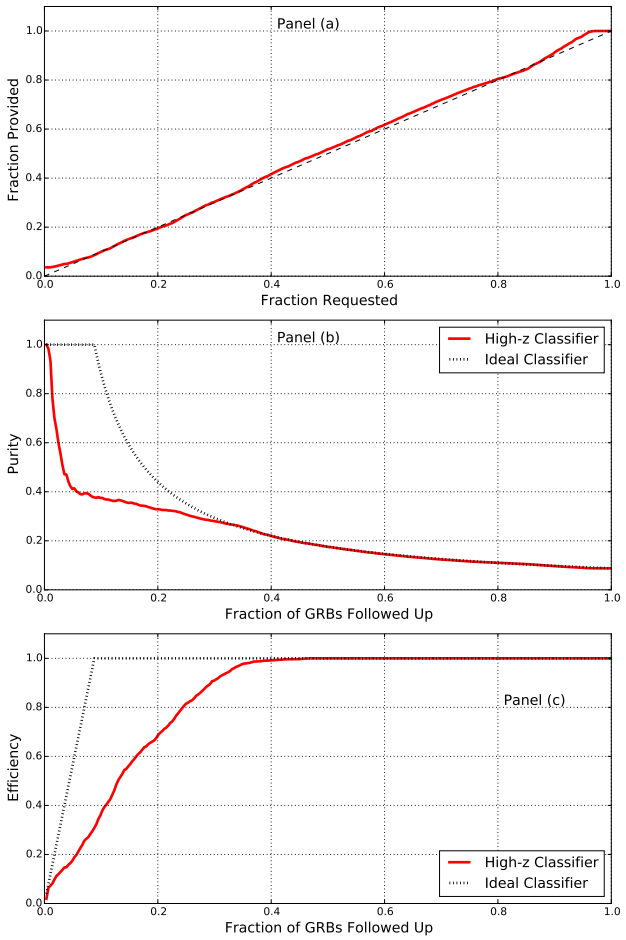

Figure 4. Performance curves for *high-z* classifier. The top panel (a) compares the fraction of bursts recommended for follow up and the fraction requested from the classifier. The middle panel (b) shows the purity of the burst sample selected for follow-up (the fraction of bursts that were followed up that are actually at high redshift). The bottom panel (c) shows the efficiently of the classifier (the fraction of all high-redshift bursts that were followed up).

search to approximately optimize input parameters of the RF regressor. For this purpose we used the "leave-one-out" cross-validation method that for N bursts consists of N runs with $N-1$ instances used for training and one for testing. We leverage the stochastic nature of RF training to increase the signal-to-noise ratio of the final cross-validated performance estimate by repeating the process 10 times with different seeds. The quality of prediction is measured using the Pearson correlation coefficient between *machine-z* output and true redshift. Table 2 defines the search grid for approximate parameter optimization. In this case we found that a forrest of 100 fully developed trees (with as little as one burst per leaf node) and $m = 5$ random features per split provides the best results. This is different from parameters adopted in section 3.2. The resulting correlation coefficient is 0.52.

4.2 Regression Feature Importance

Feature importance for *machine-z* is determined using a method similar to Section 3.3 except for the objective function. The area under the ROC curve is now replaced by the redshift correlation coefficient. Until the number of selected input features reaches $m = 5$ all features are randomized during node splitting. The relative importance of various features is shown in Fig. 5. The correlation coefficient starts from a sub-optimal value for the first feature, then increases, eventually flattens after the 11-th feature, and then slowly decreases beyond 16-th feature. We selected the first 11 features in this plot as input features for the *machine-z* estimator. Adding features beyond 11 does not improve predictions and increases the risk of overfitting or diluting the signal with noisy features.

4.3 Correction for Noise and Imbalance

A comparison between *machine-z* predictions and true redshift for GRBs in the training set is presented in Fig. 6. While there is a good correlation between the predicted and the actual redshift, the range of the output values is squeezed relative to the input. This appears to be a consequence of the interaction between noisy features and the fact that high-redshift bursts are strongly underrepresented in training data (see Fig. 1). When high-redshift bursts are given higher weights (e.g. by including multiple copies of the same burst in training data), the bias observed in Fig. 6 changes. A thorough investigation of this behavior is beyond the scope of this paper and will be presented elsewhere. For this initial release of the algorithm we introduce a simple linear correction that shifts and stretches the range of the output redshift values while preserving the correlation coefficient. The final corrected redshift predictions are computed using a straight line fit to data in Fig. 6 and taking into account the cross-validation uncertainty in *machine-z* output: $z_{\rm corrected} = (z_{\rm uncorrected} - 1.07 \pm 0.05)/(0.35 \pm 0.03)$.

The final corrected *machine-z* predictions as a function of the true redshift are shown in Fig. 7. The range of the output is now similar to that of the input and the correlation coefficient is the same as in Figure 6.

There are several interesting trends to note in Fig. 7. First, the lower right area of the plot is not populated. This means that in most cases *machine-z* does not fail to recognize a high-redshift burst. Second, the density of bursts peaks roughly along the dashed line, so for a significant fraction of bursts the redshift estimate is close to the true value. This can be seen more clearly in Fig. 8 showing the distribution of the relative differences between *machine-z* estimates and actual redshifts. Third, the algorithm does occasionally predict a high redshift for a low-z burst as shown by the upper left portion of the plot. Even though following up these false positives will tend to waste some telescope time, *machine-z* will rarely miss the all important high-z bursts. An inconvenient side effect of our simple correction for the redshift bias is that for a few GRBs the predicted redshift is negative. This is not a problem as long as the tool is used to select high-redshift GRBs, as the negative predictions only occurr at low redshift.

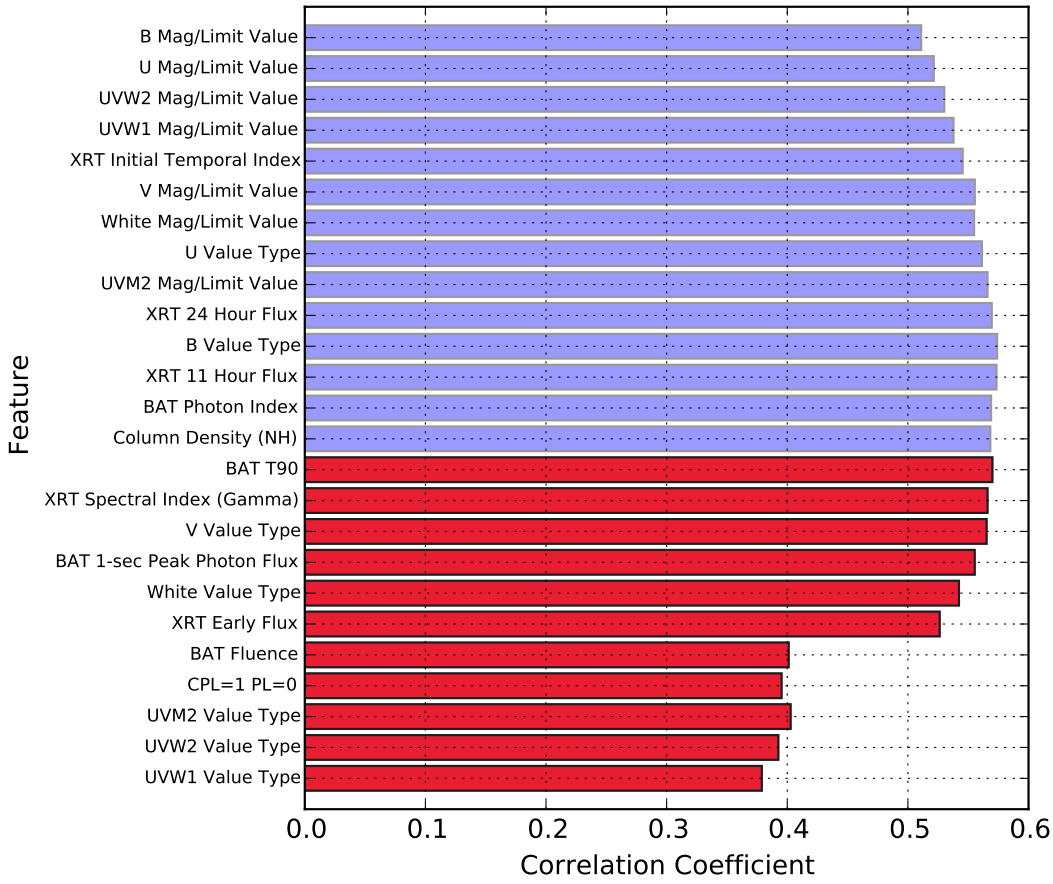

Figure 5. Relative importance of *machine-z* regression features. The Pearson correlation coefficient is shown as a function of the next best feature starting from the single best feature at the bottom of the plot. Features selected for the final *machine-z* estimation are shown in red.

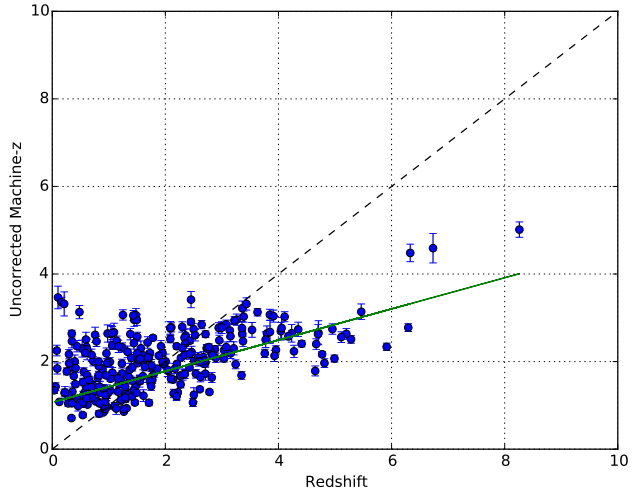

Figure 6. Comparison of uncorrected *machine-z* predictions with true redshift. The correlation coefficient between the two quantities is 0.57. The best straight line fit is show in green.

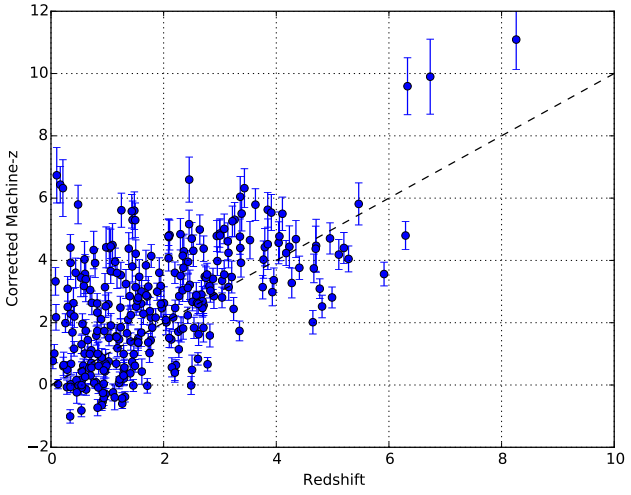

Figure 7. Comparison of corrected *machine-z* predictions with true redshift. The correlation coefficient between the two quantities is 0.57.

5 DISCUSSION

5.1 Comparison with Previous Work

Morgan et al. (2012) was the first to apply machine learned

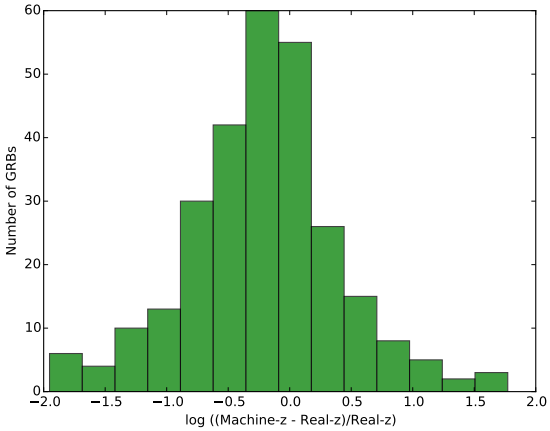

Figure 8. Distribution of relative differences between *machine-z* predictions and true redshifts.

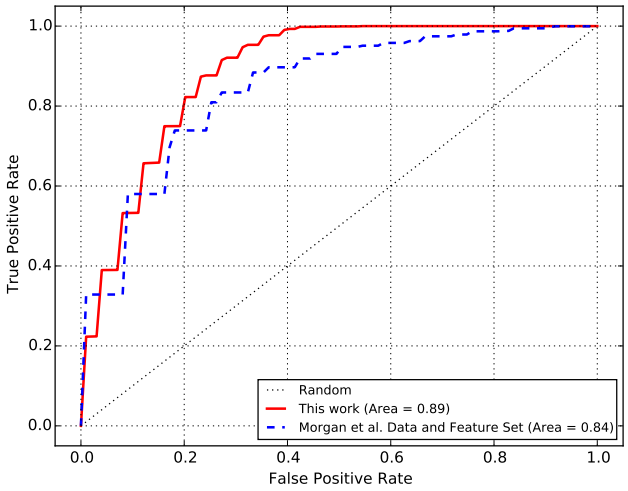

Figure 9. Comparison of ROC curves for *high-z* classifier based on our input data and feature set with Morgan et al. data and feature set. The area under the ROC curve is 0.89 for this work and 0.84 for Morgan et al..

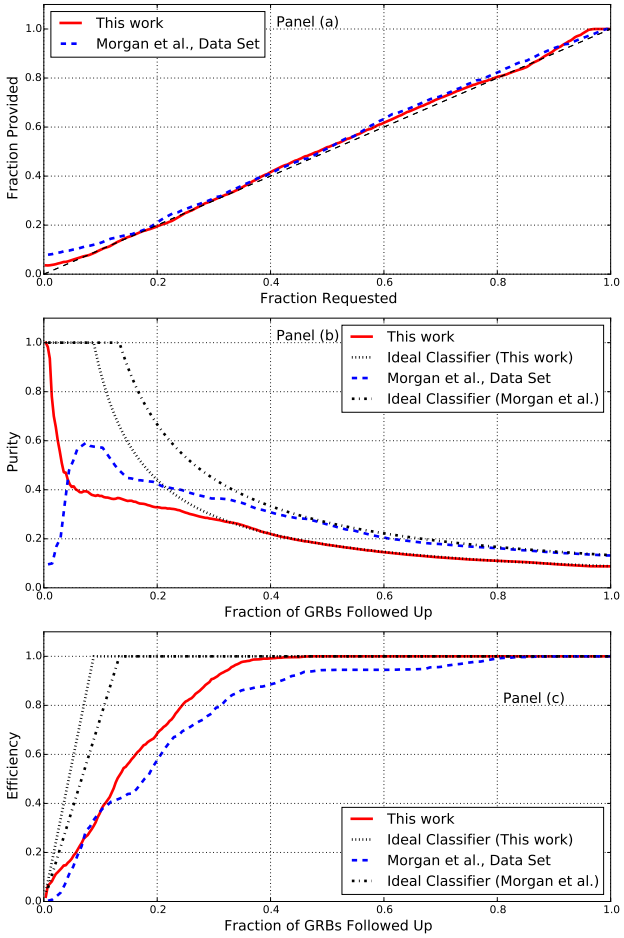

Figure 10. Performance of *high-z* classification using our input data and feature set compared to results based on the input data and feature set of Morgan et al. (2012). The content of each panel is analogous to Figure 4.

classification to screen high redshift GRBs using promptly available *Swift* data. We compared our GRB sample and feature set with the Morgan et al. (2012) data available in machine readable form. Small differences in the RF implementation between those two studies have no bearing on this comparison. The ROC curves for the two data sets are given in Fig. 9. The ROC curve corresponding to our data set (red curve) has a slightly larger area than that for the Morgan et al. (2012) data (blue curve). Note that the red curve rises to 100% recall more rapidly than the blue curve.

Fig. 10 presents another performance comparison of the two data sets. As shown by the top panel (a), there is no significant difference in the fraction of bursts recommended for follow-up versus the requested fraction. However, there is a significant difference in the purity of the burst sample selected for follow-up shown in the middle panel (b). Our *high-z* classifier returns samples of very high purity when the fraction of bursts that can be observed is low. In contrast with that, the Morgan et al. (2012) data set starts near zero purity at low followup fractions and peaks at $\approx 60\%$ around $F = 10\%$. Furthermore, the purity curve delivered by *high-z* is similar in shape to the ideal purity curve and merges with the ideal curve around the follow-up fraction $F \approx 35\%$. By comparison, the Morgan et al. (2012) data set shows a qualitatively different shape and an undesirable very low purity at small follow-up fractions. Finally, the bottom panel (c) in Fig. 10 compares efficiency curves of the two data sets. Both curves show the expected gradual rise from zero. The *high-z* data set presented here reaches the curve corresponding to a perfect classifier around follow-up fraction $F \approx 40\%$. The Morgan et al. (2012) data set does not reach the ideal curve until $\approx 80\%$.

This difference in behavior may be explained by a different approach to integrating XRT and UVOT measurements with non-detections. While the Morgan et al. (2012) classifier uses 12 features, our *high-z* algorithm uses 8 features. Some features such as the XRT column density are common to both data sets. However, Morgan et al. (2012) use only one measurement from XRT and only one feature from UVOT. The latter is limited to a yes or no flag indicating the existence of a UVOT detection in any photometric band.

Table 3. Validation sample: GRBs discovered in 2015 by *Swift* with redshift measurements. Gray rows mark candidate high-z bursts selected by both *high-z* classifier and *machine-z* estimator.

GRB	Redshift	*Machine-z*	High-z?	Score (Q value)
151112A	4.1	5.47 ± 0.56	True	<17.1 %
151111A	3.5	4.81 ± 0.59	True	<7.5 %
151031A	1.17	3.05 ± 0.39	True	<22.8 %
151029A	1.42	6.1 ± 0.71	False	<64.8 %
151027B	4.06	5.2 ± 0.59	True	<7.3 %
151027A	0.81	0.32 ± 0.21	False	<91.5 %
151021A	2.33	3.69 ± 0.38	True	<12.4 %
150915A	1.97	1.38 ± 0.39	False	<55.3 %
150910A	1.36	2.7 ± 0.41	False	<55.7 %
150821A	0.76	1.37 ± 0.34	True	<12.3 %
150818A	0.28	0.34 ± 0.2	False	<56.2 %
150727A	0.31	1.13 ± 0.3	False	<40.2 %
150424A	<3.0	-0.28 ± 0.18	False	<47.3 %
150423A	1.39	-0.26 ± 0.21	True	<4.5 %
150413A	3.2	0.75 ± 0.29	False	<36.8 %
150403A	2.06	2.36 ± 0.33	False	<62.9 %
150323A	0.59	3.81 ± 0.38	True	<29.1 %
150314A	1.76	2.99 ± 0.43	False	<31.8 %
150301B	1.52	6.06 ± 0.78	False	<39.8 %
150206A	2.09	2.81 ± 0.39	False	<35.3 %
150120A	0.46	0.81 ± 0.26	False	<74.2 %
150101B	0.09	0.71 ± 0.35	False	<73.9 %

Our approach, by contrast, is to use multiple features from XRT and utilize all available information on detections and non-detections across all UVOT bands. As one can see from Fig. 2, both detections and non-detections in various UVOT bands play an important role in the *high-z* classifier.

5.2 Validation

The training data for our *high-z* classifier and *machine-z* estimator is limited to GRBs discovered by *Swift* prior to 2015 for which a redshift measurement is available. In 2015 *Swift* found 22 additional bursts that have a spectroscopic redshift. We used the 2015 sample as a validation set to investigate the effectiveness of our redshift prediction algorithms. The results for individual bursts in the test sample are shown in Table 3. Two out of 22 bursts (GRB 151112A and GRB 151027B) qualify as high-redshift according to our classification in section 2.1 ($z > 4$). Both algorithms flag them as having high redshift. The Q scores for these two bursts also indicate that follow-up is recommended if the requested follow-up fraction is at least 20% of all GRBs. In addition to these two clear cut cases, our classifier identified 8 other bursts in the validation sample as high-z. However, out of those 8 bursts only one (GRB 151111A) has a *machine-z* estimate $z > 4$. There are also two low-z bursts (GRB 151029A and GRB 150301B) with predicted redshifts above 4. These outcomes are consistent with our performance estimates and confirm the usefulness of our tools in prioritizing follow-up observations of candidate high-redshift GRBs. We can expect that the most robust results will be obtained if both *high-z* classifier and *machine-z* estimator predict high redshift. In this case we would have selected three candidate high-z bursts shown as gray rows in Table 3, two real ones and a single false positive. Note that the false positive (GRB 151111A) is a burst with an intermediate redshift $z = 3.5$.

5.3 Extensions to Other Data Sources

The present paper addresses redshift prediction for *Swift* GRBs. Transfering a trained classifier from one data set to another is very important, but typically challenging. Unfortunately, in most cases the performance is strongly degraded even if differences between data sources are purely incidental (e.g. slightly different energy ranges of flux measurements or different estimators of model parameters). If the new features are qualitatively similar to the old ones, one can shift and rescale the numbers to approximately match the distribution of each new and old feature. This requires only a modest amount of new data and may be a productive approach for future missions similar to *Swift* including Space-based multi-band astronomical Variable Objects Monitor (SVOM; Cordier et al. (2015)). In other cases we are forced to build new training sets and that can be time consuming.

Another possible approach would be to consider lower level data such as time-resolved spectra of prompt GRB emission. Generic intermediate level features can be obtained for example from wavelet analysis that captures the intrinsic structure of the data and then apply a high level classfier such as RF (Ukwatta and Wozniak 2016). Those "abstract" features may prove more transferable from one data set to another and may eventually facilitate early redshift prediction for GRBs detected by future missions such as SVOM.

6 SUMMARY

We presented a method for selecting candidate high redshift gamma-ray bursts that can be used to prioritize follow-up observations. The algorithm utilizes numerical and categorical features from all three instruments onboard the *Swift* satellite that are readily available within the first few hours after a GRB discovery. We independently developed *high-z* classification and *machine-z* regression tools based on algorithms and features tailored to each task. A subset of features that provides most information to support redshift prediction was identified in both cases. The results were validated using a small sample of recently discovered GRBs that were not included in training data. The most robust selection of high redshift bursts is achieved by combining the classification and regression output to make the final prediction.

ACKNOWLEDGMENTS

This work was funded by the US Department of Energy. We acknowledge support from the Laboratory Directed Research and Development program at Los Alamos National Laboratory. We also thank the anonymous referee for comments that significantly improved the paper.

REFERENCES

L. Breiman. Random Forests. *Machine Learning*, 45:5–32, January 2001.

S. Campana, G. Tagliaferri, D. Malesani, L. Stella, P. D'Avanzo, G. Chincarini, and S. Covino. Near real-time selection of high redshift GRBs with Swift. A.& A., 464: L25–L27, March 2007. doi: 10.1051/0004-6361:20066592.

S. Carliles, T. Budavári, S. Heinis, C. Priebe, and A. S. Szalay. Random Forests for Photometric Redshifts. ApJ., 712:511–515, March 2010. doi: 10.1088/0004-637X/712/1/511.

R. Caruana and A. Niculescu-Mizil. An Empirical Comparison of Supervised Learning Algorithms. *Proceedings of the 23rd International Conference on Machine Learning, Pittsburgh, PA*, January 2006.

B. Cordier, J. Wei, J.-L. Atteia, S. Basa, A. Claret, F. Daigne, J. Deng, Y. Dong, O. Godet, A. Goldwurm, D. Götz, X. Han, A. Klotz, C. Lachaud, J. Osborne, Y. Qiu, S. Schanne, B. Wu, J. Wang, C Wu, L. Xin, B. Zhang, and S.-N. Zhang. The SVOM gamma-ray burst mission. *ArXiv:1512.03323*, December 2015.

A. Cucchiara, A. J. Levan, D. B. Fox, N. R. Tanvir, T. N. Ukwatta, E. Berger, T. Krühler, A. Küpcü Yoldaş, X. F. Wu, K. Toma, J. Greiner, F. E. Olivares, A. Rowlinson, L. Amati, T. Sakamoto, K. Roth, A. Stephens, A. Fritz, J. P. U. Fynbo, J. Hjorth, D. Malesani, P. Jakobsson, K. Wiersema, P. T. O'Brien, A. M. Soderberg, R. J. Foley, A. S. Fruchter, J. Rhoads, R. E. Rutledge, B. P. Schmidt, M. A. Dopita, P. Podsiadlowski, R. Willingale, C. Wolf, S. R. Kulkarni, and P. D'Avanzo. A Photometric Redshift of z $\sim$ 9.4 for GRB 090429B. ApJ., 736:7, July 2011. doi: 10.1088/0004-637X/736/1/7.

G. Cusumano, V. Mangano, G. Chincarini, A. Panaitescu, D. N. Burrows, V. La Parola, T. Sakamoto, S. Campana, T. Mineo, G. Tagliaferri, L. Angelini, S. D. Barthelemy, A. P. Beardmore, P. T. Boyd, L. R. Cominsky, C. Gronwall, E. E. Fenimore, N. Gehrels, P. Giommi, M. Goad, K. Hurley, J. A. Kennea, K. O. Mason, F. Marshall, P. Mészáros, J. A. Nousek, J. P. Osborne, D. M. Palmer, P. W. A. Roming, A. Wells, N. E. White, and B. Zhang. Gamma-ray bursts: Huge explosion in the early Universe. Nature, 440:164, March 2006. doi: 10.1038/440164a.

A. D'Isanto, S. Cavuoti, M. Brescia, C. Donalek, G. Longo, G. Riccio, and S. G. Djorgovski. An analysis of feature relevance in the classification of astronomical transients with machine learning methods. *ArXiv:1601.03931*, January 2016.

L. du Buisson, N. Sivanandam, B. A. Bassett, and M. Smith. Machine learning classification of SDSS transient survey images. MNRAS, 454:2026–2038, December 2015. doi: 10.1093/mnras/stv2041.

S. A. Farrell, T. Murphy, and K. K. Lo. Autoclassification of the Variable 3XMM Sources Using the Random Forest Machine Learning Algorithm. ApJ., 813:28, November 2015. doi: 10.1088/0004-637X/813/1/28.

D. A. Frail, P. B. Cameron, M. Kasliwal, E. Nakar, P. A. Price, E. Berger, A. Gal-Yam, S. R. Kulkarni, D. B. Fox, A. M. Soderberg, B. P. Schmidt, E. Ofek, and S. B. Cenko. An Energetic Afterglow from a Distant Stellar Explosion. ApJ. Lett., 646:L99–L102, August 2006. doi: 10.1086/506934.

N. Gehrels, G. Chincarini, P. Giommi, K. O. Mason, J. A. Nousek, A. A. Wells, N. E. White, S. D. Barthelmy, D. N. Burrows, L. R. Cominsky, K. C. Hurley, F. E. Marshall, P. Mészáros, P. W. A. Roming, L. Angelini, L. M. Barbier, T. Belloni, S. Campana, P. A. Caraveo, M. M. Chester, O. Citterio, T. L. Cline, M. S. Cropper, J. R. Cummings, A. J. Dean, E. D. Feigelson, E. E. Fenimore, D. A. Frail, A. S. Fruchter, G. P. Garmire, K. Gendreau, G. Ghisellini, J. Greiner, J. E. Hill, S. D. Hunsberger, H. A. Krimm, S. R. Kulkarni, P. Kumar, F. Lebrun, N. M. Lloyd-Ronning, C. B. Markwardt, B. J. Mattson, R. F. Mushotzky, J. P. Norris, J. Osborne, B. Paczynski, D. M. Palmer, H.-S. Park, A. M. Parsons, J. Paul, M. J. Rees, C. S. Reynolds, J. E. Rhoads, T. P. Sasseen, B. E. Schaefer, A. T. Short, A. P. Smale, I. A. Smith, L. Stella, G. Tagliaferri, T. Takahashi, M. Tashiro, L. K. Townsley, J. Tueller, M. J. L. Turner, M. Vietri, W. Voges, M. J. Ward, R. Willingale, F. M. Zerbi, and W. W. Zhang. The Swift Gamma-Ray Burst Mission. ApJ., 611:1005–1020, August 2004. doi: 10.1086/422091.

D. A. Goldstein, C. B. D'Andrea, J. A. Fischer, R. J. Foley, R. R. Gupta, R. Kessler, A. G. Kim, R. C. Nichol, P. E. Nugent, A. Papadopoulos, M. Sako, M. Smith, M. Sullivan, R. C. Thomas, W. Wester, R. C. Wolf, F. B. Abdalla, M. Banerji, A. Benoit-Lévy, E. Bertin, D. Brooks, A. Carnero Rosell, F. J. Castander, L. N. da Costa, R. Covarrubias, D. L. DePoy, S. Desai, H. T. Diehl, P. Doel, T. F. Eifler, A. Fausti Neto, D. A. Finley, B. Flaugher, P. Fosalba, J. Frieman, D. Gerdes, D. Gruen, R. A. Gruendl, D. James, K. Kuehn, N. Kuropatkin, O. Lahav, T. S. Li, M. A. G. Maia, M. Makler, M. March, J. L. Marshall, P. Martini, K. W. Merritt, R. Miquel, B. Nord, R. Ogando, A. A. Plazas, A. K. Romer, A. Roodman, E. Sanchez, V. Scarpine, M. Schubnell, I. Sevilla-Noarbe, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, M. E. C. Swanson, G. Tarle, J. Thaler, and A. R. Walker. Automated Transient Identification in the Dark Energy Survey. AJ., 150:82, September 2015. doi: 10.1088/0004-6256/150/3/82.

K. Ioka. The Cosmic Dispersion Measure from Gamma-Ray Burst Afterglows: Probing the Reionization History and the Burst Environment. ApJ. Lett., 598:L79–L82, December 2003. doi: 10.1086/380598.

N. Kawai, G. Kosugi, K. Aoki, T. Yamada, T. Totani, K. Ohta, M. Iye, T. Hattori, W. Aoki, H. Furusawa, K. Hurley, K. S. Kawabata, N. Kobayashi, Y. Komiyama, Y. Mizumoto, K. Nomoto, J. Noumaru, R. Ogasawara, R. Sato, K. Sekiguchi, Y. Shirasaki, M. Suzuki, T. Takata, T. Tamagawa, H. Terada, J. Watanabe, Y. Yatsu, and A. Yoshida. An optical spectrum of the afterglow of a γ-ray burst at a redshift of z = 6.295. Nature, 440:184–186, March 2006. doi: 10.1038/nature04498.

C. Koen. Multiple regression of GRB luminosity on light-curve properties. MNRAS, 396:1499–1506, July 2009. doi: 10.1111/j.1365-2966.2009.14795.x.

C. Koen. Predicting gamma-ray burster redshifts from their prompt emission properties. MNRAS, 401:1369–1374, January 2010. doi: 10.1111/j.1365-2966.2009.15737.x.

D. Q. Lamb and D. E. Reichart. Gamma-Ray Bursts as a Probe of the Very High Redshift Universe. ApJ., 536: 1–18, June 2000. doi: 10.1086/308918.

R. A. Mesler, D. J. Whalen, J. Smidt, C. L. Fryer, N. M. Lloyd-Ronning, and Y. M. Pihlström. The First Gamma-Ray Bursts in the Universe. ApJ., 787:91, May 2014. doi: 10.1088/0004-637X/787/1/91.

A. N. Morgan, J. Long, J. W. Richards, T. Broderick, N. R. Butler, and J. S. Bloom. Rapid, Machine-learned Resource Allocation: Application to High-redshift Gamma-Ray Burst Follow-up. ApJ., 746:170, February 2012. doi: 10.1088/0004-637X/746/2/170.

F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay. Scikit-learn: Machine Learning in Python. *ArXiv:1201.0490*, January 2012.

R. Salvaterra, S. Campana, G. Chincarini, G. Tagliaferri, and S. Covino. On the detection of very high redshift gamma-ray bursts with Swift. MNRAS, 380:L45–L48, September 2007. doi: 10.1111/j.1745-3933.2007.00345.x.

R. Salvaterra, M. Della Valle, S. Campana, G. Chincarini, S. Covino, P. D'Avanzo, A. Fernández-Soto, C. Guidorzi, F. Mannucci, R. Margutti, C. C. Thöne, L. A. Antonelli, S. D. Barthelmy, M. de Pasquale, V. D'Elia, F. Fiore, D. Fugazza, L. K. Hunt, E. Maiorano, S. Marinoni, F. E. Marshall, E. Molinari, J. Nousek, E. Pian, J. L. Racusin, L. Stella, L. Amati, G. Andreuzzi, G. Cusumano, E. E. Fenimore, P. Ferrero, P. Giommi, D. Guetta, S. T. Holland, K. Hurley, G. L. Israel, J. Mao, C. B. Markwardt, N. Masetti, C. Pagani, E. Palazzi, D. M. Palmer, S. Piranomonte, G. Tagliaferri, and V. Testa. GRB090423 at a redshift of z $\sim$ 8.1. Nature, 461:1258–1260, October 2009. doi: 10.1038/nature08445.

N. R. Tanvir. The highest redshift gamma-ray bursts. *ArXiv:1307.6156*, July 2013.

N. R. Tanvir, D. B. Fox, A. J. Levan, E. Berger, K. Wiersema, J. P. U. Fynbo, A. Cucchiara, T. Krühler, N. Gehrels, J. S. Bloom, J. Greiner, P. A. Evans, E. Rol, F. Olivares, J. Hjorth, P. Jakobsson, J. Farihi, R. Willingale, R. L. C. Starling, S. B. Cenko, D. Perley, J. R. Maund, J. Duke, R. A. M. J. Wijers, A. J. Adamson, A. Allan, M. N. Bremer, D. N. Burrows, A. J. Castro-Tirado, B. Cavanagh, A. de Ugarte Postigo, M. A. Dopita, T. A. Fatkhullin, A. S. Fruchter, R. J. Foley, J. Gorosabel, J. Kennea, T. Kerr, S. Klose, H. A. Krimm, V. N. Komarova, S. R. Kulkarni, A. S. Moskvitin, C. G. Mundell, T. Naylor, K. Page, B. E. Penprase, M. Perri, P. Podsiadlowski, K. Roth, R. E. Rutledge, T. Sakamoto, P. Schady, B. P. Schmidt, A. M. Soderberg, J. Sollerman, A. W. Stephens, G. Stratta, T. N. Ukwatta, D. Watson, E. Westra, T. Wold, and C. Wolf. A γ-ray burst at a redshift of z $\sim$ 8.2. Nature, 461:1254–1257, October 2009. doi: 10.1038/nature08459.

T. Totani, N. Kawai, G. Kosugi, K. Aoki, T. Yamada, M. Iye, K. Ohta, and T. Hattori. Implications for Cosmic Reionization from the Optical Afterglow Spectrum of the Gamma-Ray Burst 050904 at z = 6.3. PASJ, 58: 485–498, June 2006. doi: 10.1093/pasj/58.3.485.

T. N. Ukwatta and P. R. Wozniak. Integrating Temporal and Spectral Features of Astronomical Data Using Wavelet Analysis for Source Classification. *ArXiv:1601.05529*, January 2016.

T. N. Ukwatta, T. Sakamoto, M. Stamatikos, N. Gehrels, and K. S. Dhuga. Screening High-z GRBs with BAT Prompt Emission Properties. In M. Galassi, D. Palmer, and E. Fenimore, editors, *American Institute of Physics Conference Series*, volume 1000 of *American Institute of Physics Conference Series*, pages 166–169, May 2008. doi: 10.1063/1.2943435.

T. N. Ukwatta, T. Sakamoto, K. S. Dhuga, W. C. Parke, S. D. Barthelmy, N. Gehrels, M. Stamatikos, and J. Tueller. Investigating the Possibility of Screening High-z GRBs based on BAT Prompt Emission Properties. In C. Meegan, C. Kouveliotou, and N. Gehrels, editors, *American Institute of Physics Conference Series*, volume 1133 of *American Institute of Physics Conference Series*, pages 437–439, May 2009. doi: 10.1063/1.3155945.

F. Y. Wang, V. Bromm, T. H. Greif, A. Stacy, Z. G. Dai, A. Loeb, and K. S. Cheng. Probing Pre-galactic Metal Enrichment with High-redshift Gamma-Ray Bursts. ApJ., 760:27, November 2012. doi: 10.1088/0004-637X/760/1/27.

D. E. Wright, S. J. Smartt, K. W. Smith, P. Miller, R. Kotak, A. Rest, W. S. Burgett, K. C. Chambers, H. Flewelling, K. W. Hodapp, M. Huber, R. Jedicke, N. Kaiser, N. Metcalfe, P. A. Price, J. L. Tonry, R. J. Wainscoat, and C. Waters. Machine learning for transient discovery in Pan-STARRS1 difference imaging. MNRAS, 449:451–466, May 2015. doi: 10.1093/mnras/stv292.

This paper has been typeset from a TEX/ LATEX file prepared by the author.